Hands-on Advanced Networking Topics: BGP, BGP Hijacking, MPLS, MPLS-based VPNs, Segment Routing, and others

Jorge Crichigno, Shahrin Sharif
University of South Carolina
http://ce.sc.edu/cyberinfra
jcrichigno@cec.sc.edu, ssharif@email.sc.edu

WASTC 2021 virtual Faculty Development Weeks (vFDW)
June 17, 2021

VPN Overview

- MPLS has been used to implement so-called virtual private networks (VPNs)
- MPLS VPN is perhaps the most popular and widespread application of MPLS
- It is also referred to as “peer-to-peer” VPN model
- MPLS can be used to isolate both the resources and addressing used by the customer’s VPN from that of other users crossing the ISP’s network

MPLS VPN Model

- With MPLS VPN (peer-to-peer), the service provider routers carry the customer data across the network
- The ISP routers peer directly with the customer routers at layer 3
- A routing protocol neighborship / adjacency exists between the customer and the service provider router
- The VPN model also requires privacy or isolation between the different customers

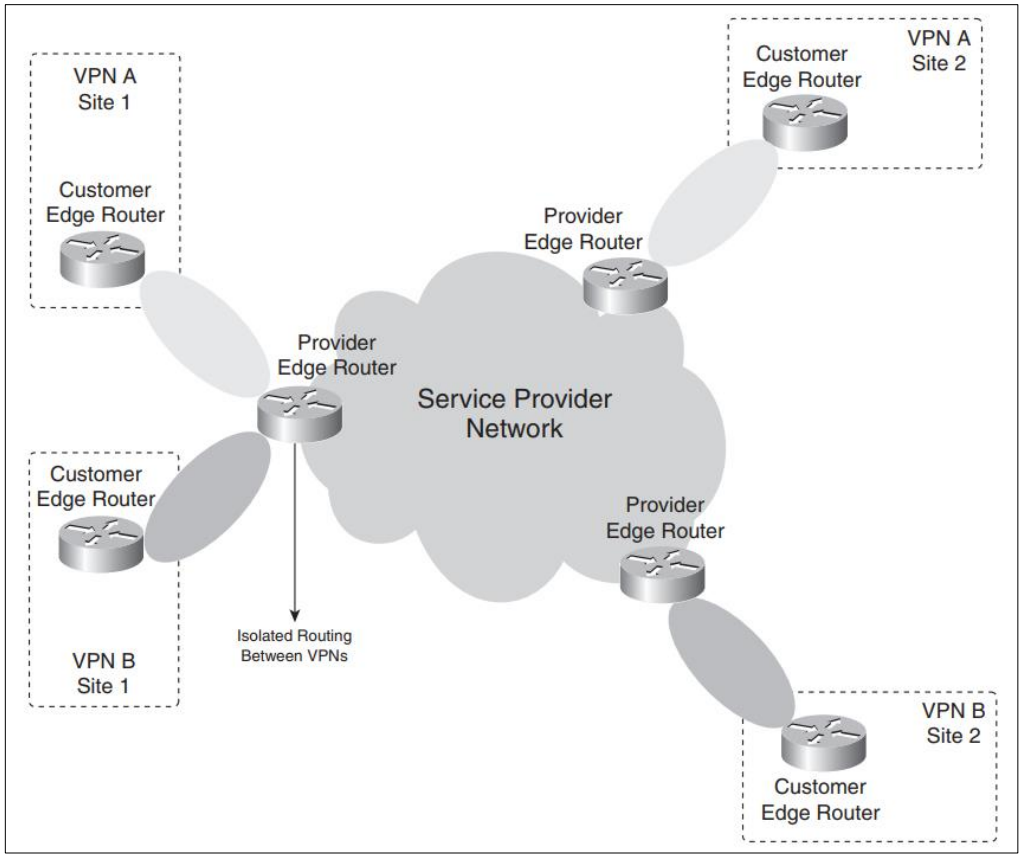

MPLS VPN Model

- Privacy is achieved by using the concept of virtual routing / forwarding (VRF) and the fact that the data is forwarded in the backbone as labeled packets
- VRFs ensure that the routing information from the different customers is kept separate
- The MPLS in the backbone ensures that the packets are forwarding based on the label information and not based on the information in the IP header

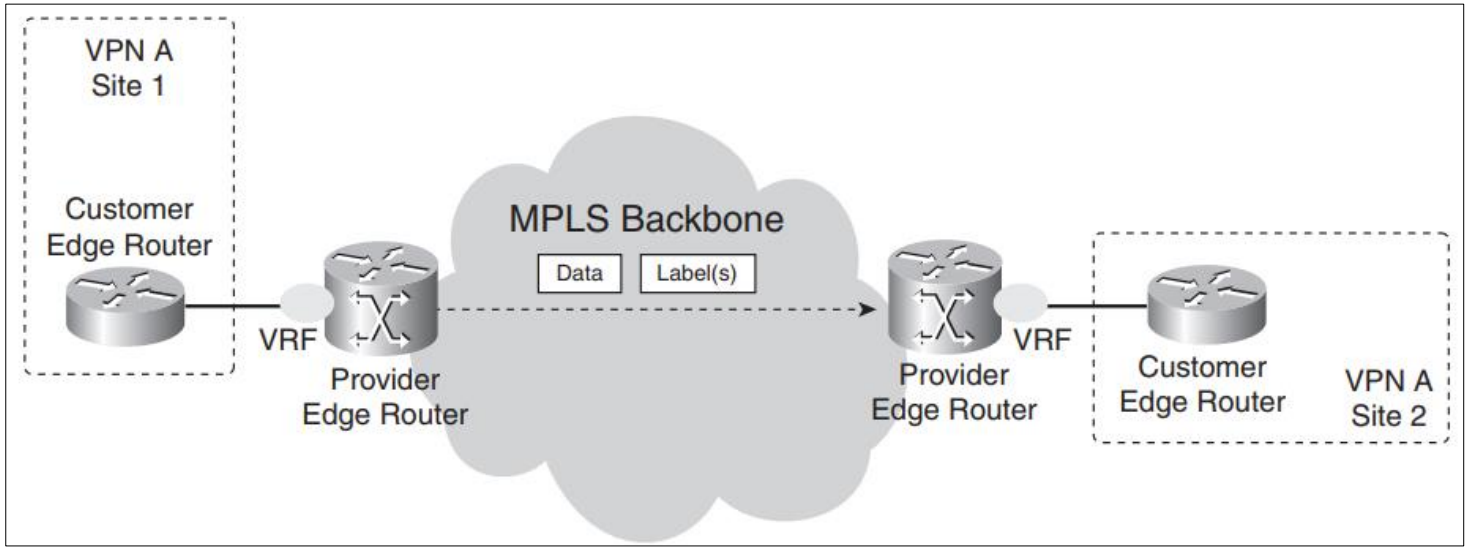

MPLS VPN Model - Terminology

- A PE router is a provider edge (PE) router. It has a connection with the customer edge (CE) at Layer 3
- A provider (P) router is a router without the direct connection to the customer routers
- P and PE routers run MPLS, distribute labels between them, forward labeled packets
- CE router has a direct layer 3 connection with the PE router; CE does not run MPLS
- A customer (C) router is a router without a direct connection with the PE router

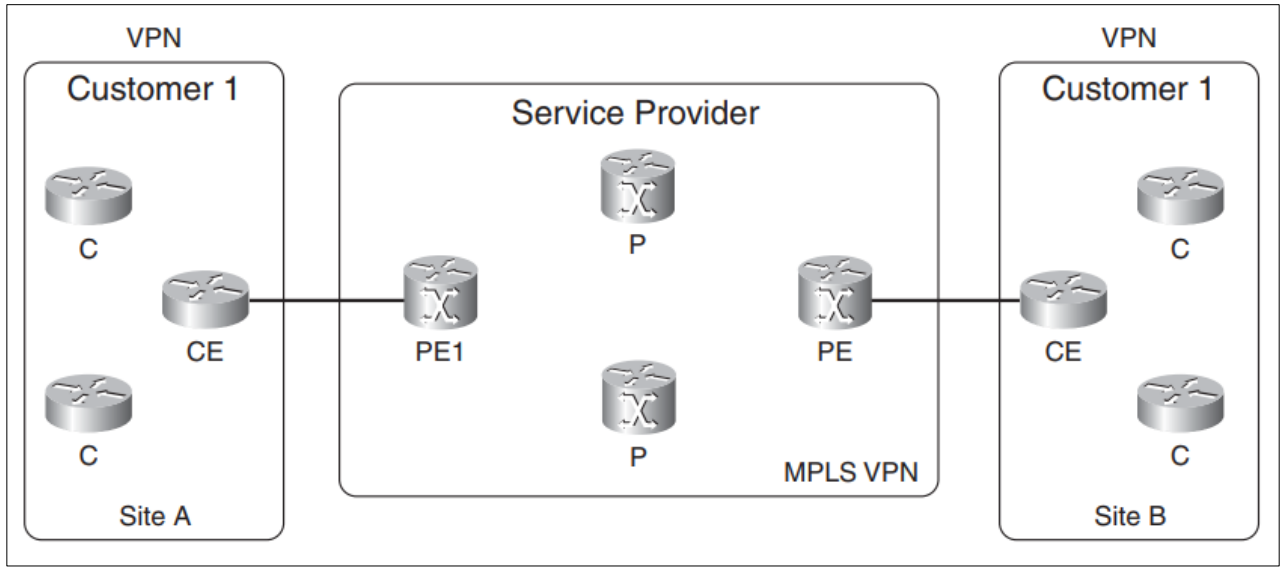

MPLS VPN Model - Terminology

- Customers of the service provider are allowed to have their own IP addressing scheme, including private IP addresses
- Different customers may even use the same (overlapping) IP addresses
- P routers are completely unaware of the VPNs
- The VPN routes are only known on the PE routers (scalable)
  - Adding a customer site means that on the PE router, only the peering with the CE router must be added
- Each VPN has its own routing table (VRF)

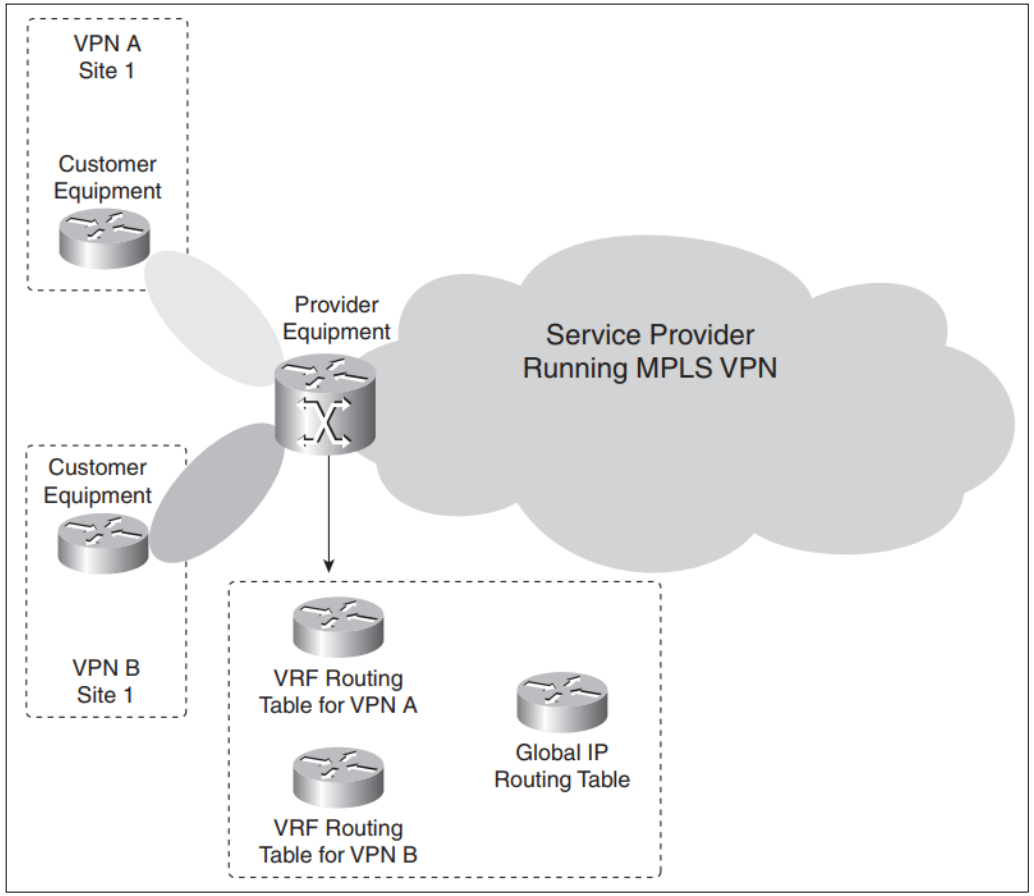

MPLS VPN Model - Terminology

- VPN prefixes are propagated across the MPLS VPN network by Multiprotocol BGP
  - Issue: when BGP carries these IPv4 prefixes across the service provider network, they must be unique
- The concept of route distinguishers (RDs) was conceived to make IPv4 prefixes unique
- Each prefix from each customer receives a unique identifier (the RD) to distinguish the same prefix from different customers
- BGP carries these prefixes between PE routers, referred to as vpnv4 prefixes

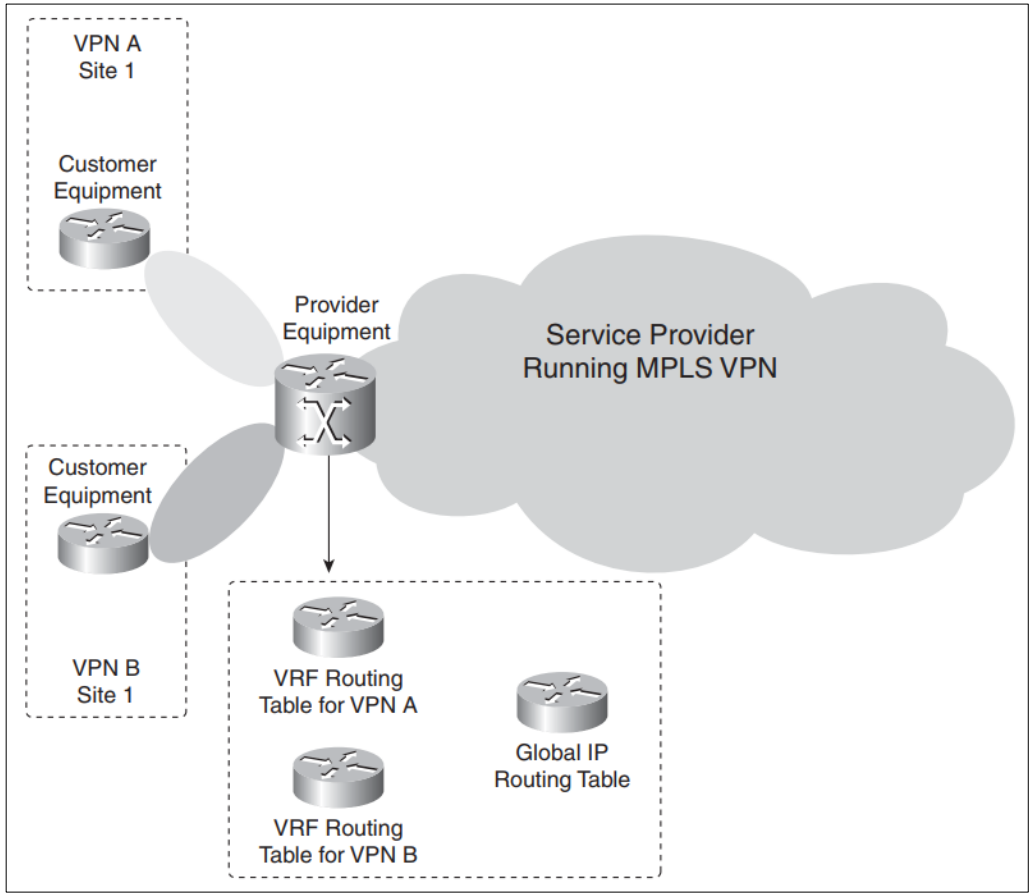

MPLS Layer 3 VPN Configuration

- PE routers run multiple VRFs for each organization (Org 1, Org 2)
- BGP advertises VPN information from PE1 to PE2

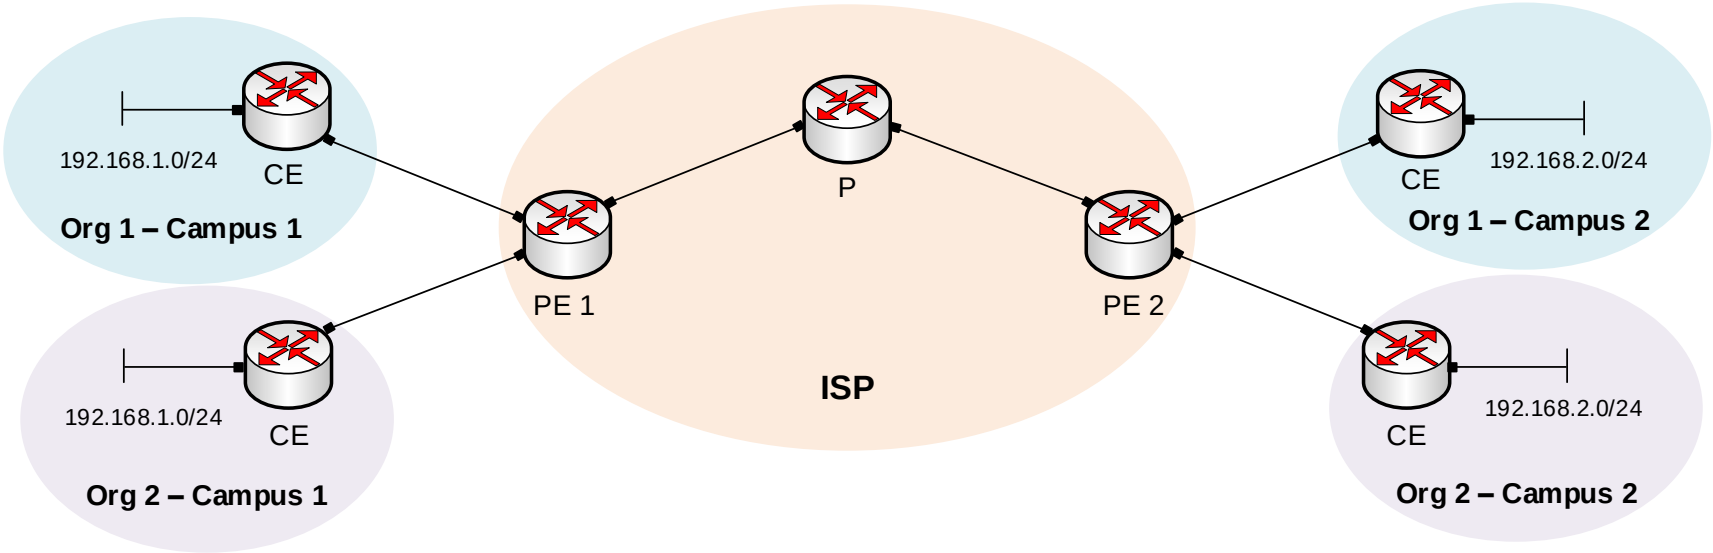

VPN Configuration for Overlapping IP Addresses

- The IPv4 prefixes carried by BGP across the ISP must be unique
- The scheme requires an identifier to distinguish between overlapping IP addresses

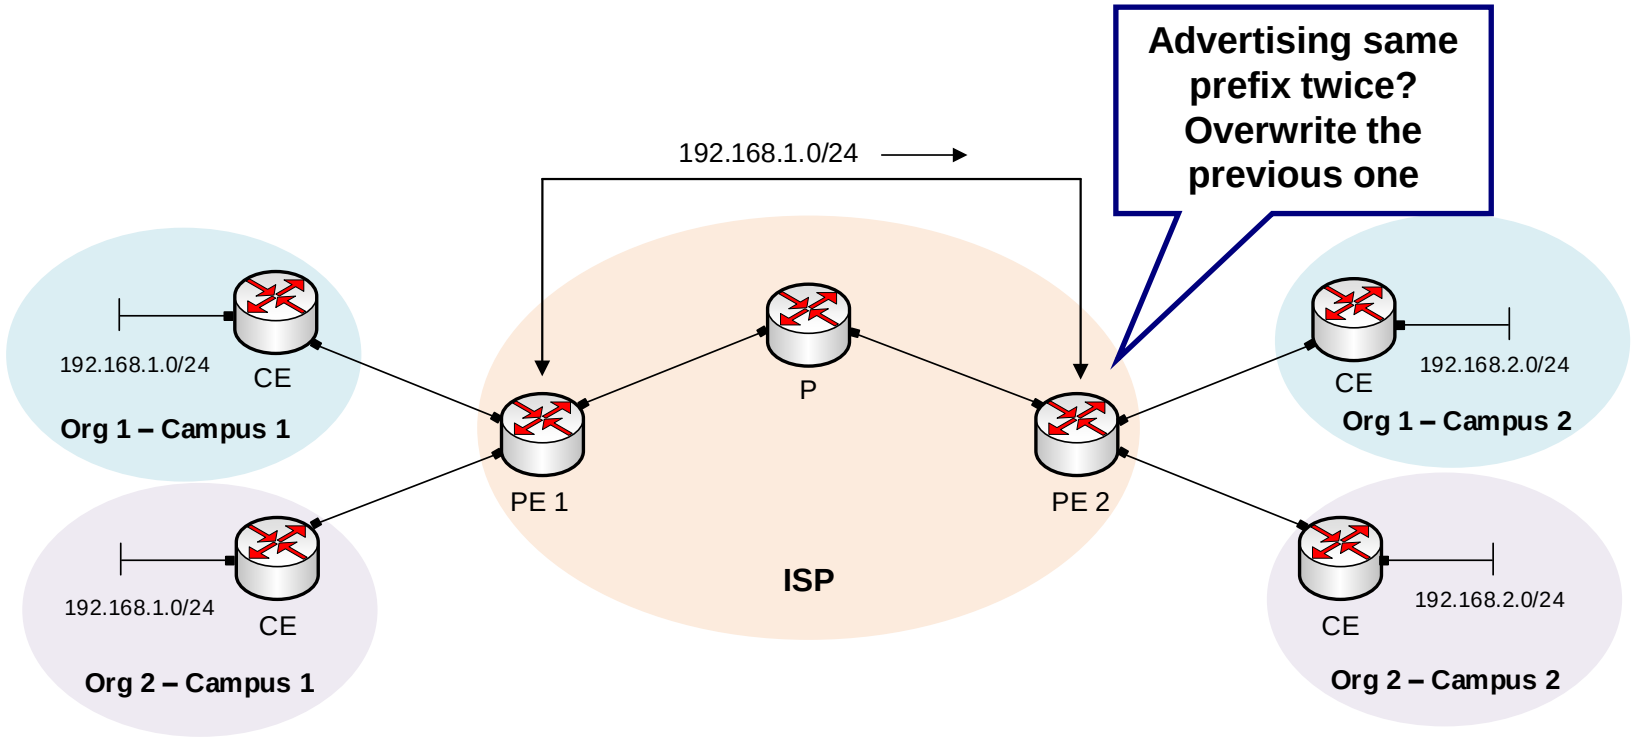

Route Distinguisher (RD)

- RD is a 64-bit field value
- RD value is added to make the VPN prefix unique

VRF	IPv4 address	Unique address
org 1	192.168.1.0/24	**1:1**:192.168.1.0/24
org 2	192.168.1.0/24	**2:2**:192.168.1.0/24
org 1	192.168.2.0/24	**1:1**:192.168.2.0/24
org 2	192.168.2.0/24	**2:2**:192.168.2.0/24

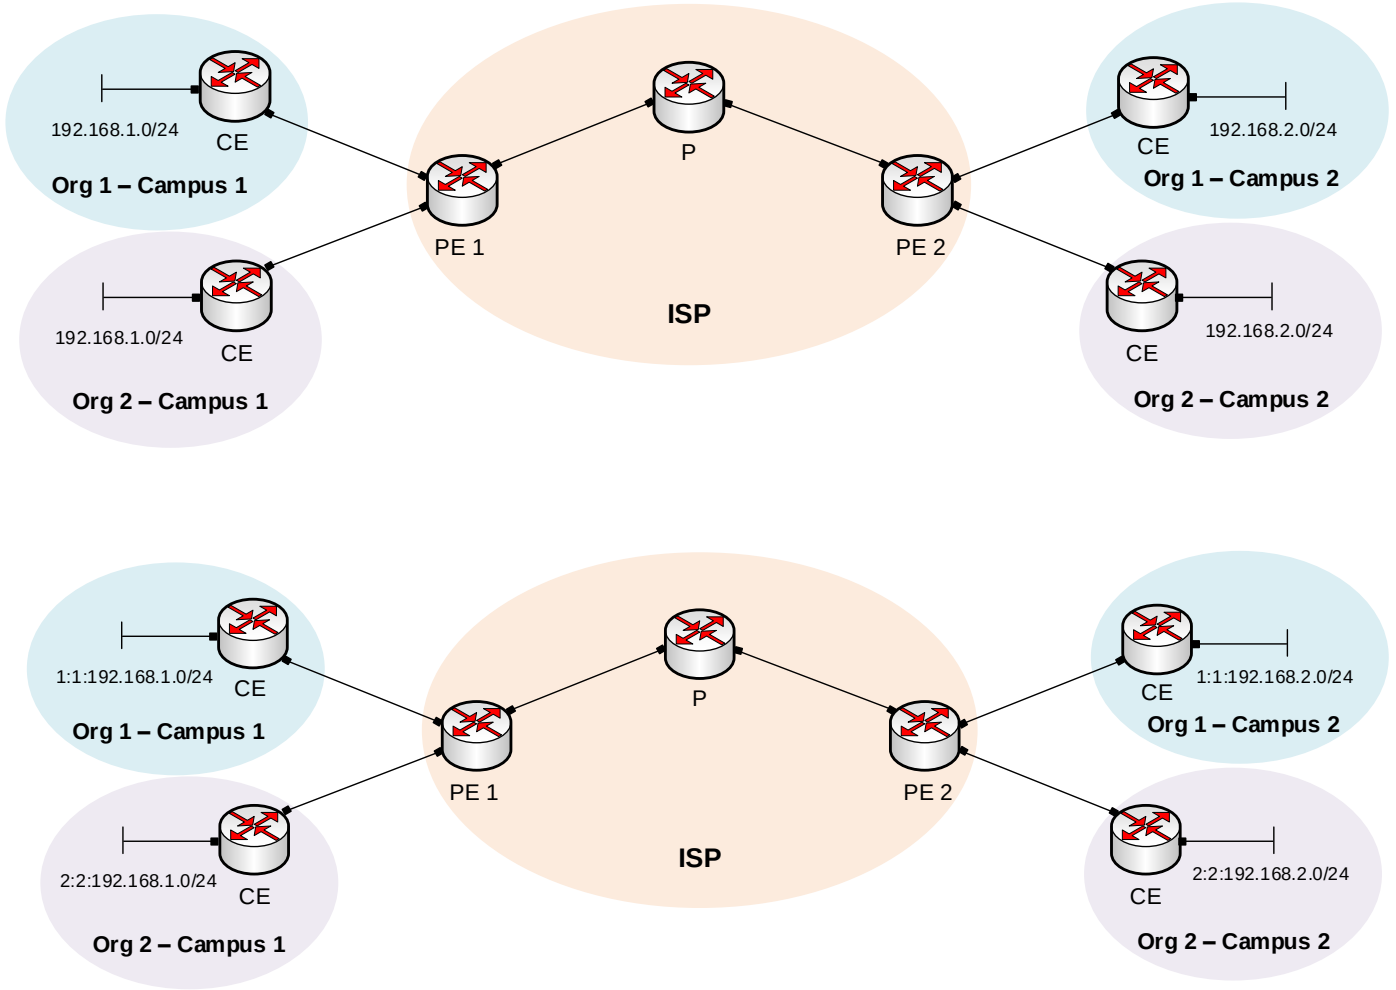

Selecting Routes to Copy to/from BGP Updates

- Selecting which vpnv4 (routes) to export from a VRF to a BGP update and which vpnv4 to import from a BGP update to a VRF are controlled by a Route Target (RT)
  - An RT indicates which routes should be imported from BGP into a VRF (routing table for a VPN)
  - Exporting an RT indicates which vpnv4 routes are redistributed from the VRF routing table into BGP
- Importing an RT means that the received vpnv4 route from BGP is checked for a matching extended community (RT) in the VRF of a router processing the BGP update
- If the result is a match, the prefix is put into the VRF routing table as an IPv4 route

Route Target (RT)

- Which routes should be imported from BGP into the VRF?
- RT is used for this purpose; it is a 64-bit value that uses the same format as RD

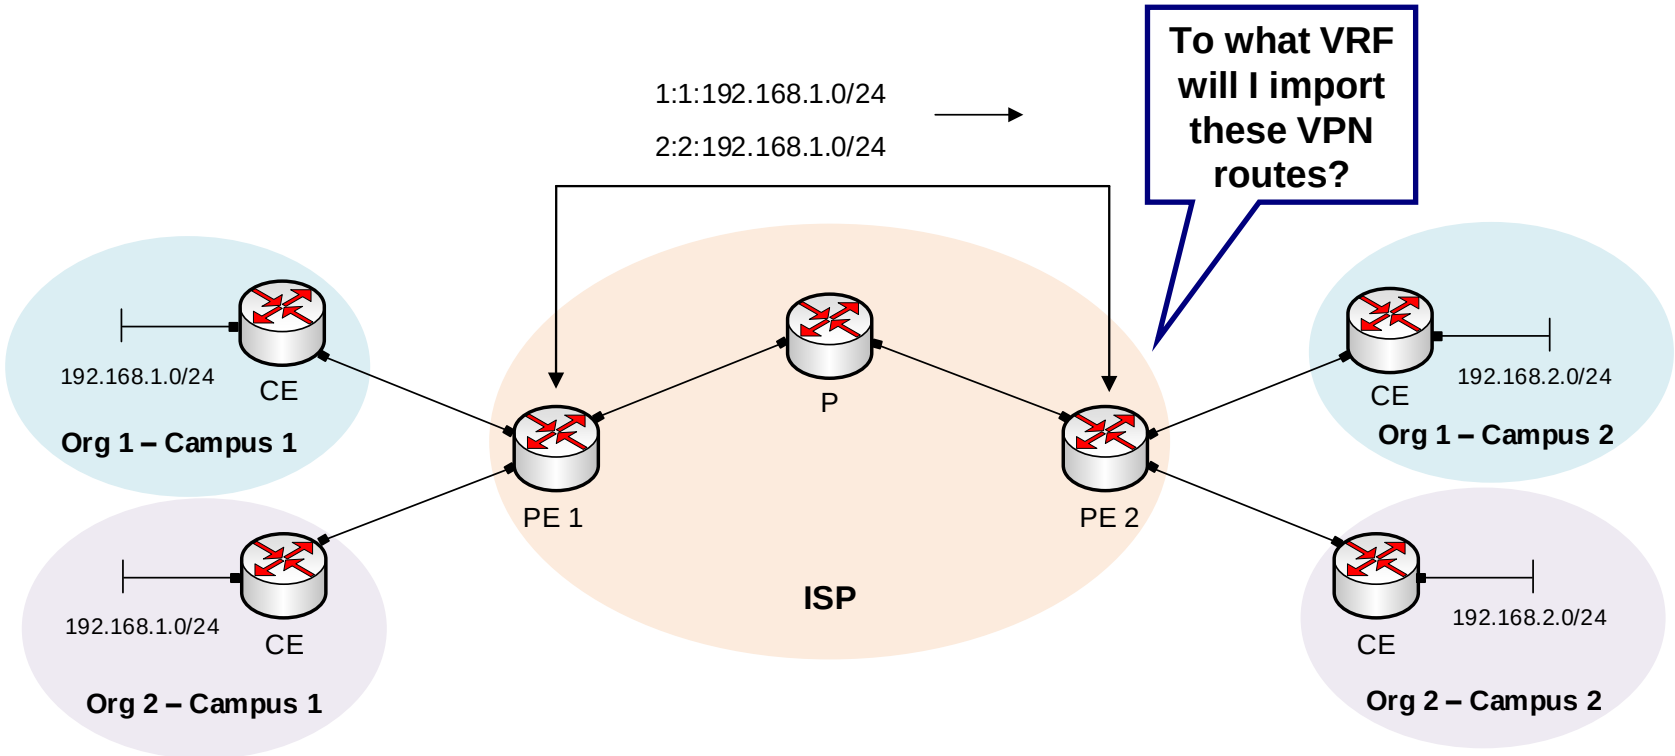

Route Target (RT)

- PE routers assign RT import and export value for each VRF
- PE1 exports RT value 100:100 for Org 1
- PE 2 selects the VRF according to the import value assigned for Org 1(a match exists between RT of BGP update and RT import configured for org1